DOCUMENT RESUME

ED 192 607 FL 011 830

AUTHOR Nahinsky, Irwin D.
TITLE Notes on a Theory of Concept Representation in Semantic Memory. Lektos: Interdisciplinary Working Papers in Language Sciences, New Series, Volume 1.
INSTITUTION Louisville Univ., Ky. Interdisciplinary Program in Linguistics.
PUB DATE 80
NOTE 47p.; Not available in paper copy due to small print size.

EDRS PRICE MF01 Plus Postage. PC Not Available from EDRS.
DESCRIPTORS Classification; *Cognitive Processes; *Concept Formation; Learning Theories; *Memory; Mnemonics; Psycholinguistics; *Semantics; Verbal Learning

ABSTRACT

A theory for representation of concepts in memory is proposed which emphasizes the association of salient exemplars with the concept. Previous theories dealing with the classification processes involved in acquiring new concepts have shown that clear category boundaries do not exist. It is proposed that present theory must account for the "fuzziness" of category boundaries found in natural languages as well as how learning is nonetheless achieved. A set of exemplars, associated with the concept through learning experience, is used for comparison with a newly encountered stimulus to ascertain if that stimulus is encoded as a set of features. Thus, the exemplars associated with the concept become a collection of intersecting feature sets. In this context, the set of features common to all exemplars is considered the set of "defining" features for the concept. A concept is considered "fuzzy" to the extent that "nondefining" features in the set of represented features are more frequent than defining features. Difficulty of verifying class membership for stimuli is hypothesized to increase with fuzziness of the concept as measured by a proposed index. Tests of the theory are suggested. (Author/PJM)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED192607

LEKTOS: INTERDISCIPLINARY WORKING PAPERS
IN LANGUAGE SCIENCES, NEW SERIES, VOL. 1

Notes on a Theory of Concept Representation in Semantic Memory

Irwin D. Nahinsky
Department of Psychology
University of Louisville

Abstract

A theory for representation of concepts in memory is proposed which emphasizes the association of salient exemplars with the concept. A set of exemplars, associated with the concept through learning experience, is used for comparison with a newly encountered stimulus to ascertain if that stimulus is an example of the concept. Each exemplar in the representation is encoded as a set of features. Thus, the exemplars associated with the concept become a collection of intersecting feature sets. In this context, the set of features common to all exemplars is considered the set of "defining" features for the concept. A concept is considered "fuzzy" to the extent that "nondefining" features in the set of represented features are more frequent than defining features. Difficulty of verifying class membership for stimuli is hypothesized to increase with fuzziness of the concept as measured by a proposed index. Tests of the theory are suggested.

The process of categorizing experience in terms of meaningful interrelationships is pervasive in human experience. We categorize all events, and we use these categorizations to make decisions. We stereotype individuals, classify objects, and categorize in some manner our daily experiences. Such processes often result in verbal labels, in which case the structure of semantic memory may incorporate a network of relationships including a mixture of abstract verbal elements, schematic representations, and even images. A number of approaches have been developed to characterize the nature of category structures (e.g., Collins & Quillian, 1969; Meyer, 1970; Rosch, 1975; Rumelhart, Lindsay, and Norman, 1972; Smith, Shoben, and Rips, 1974). The associated research has, for the most part, involved a study of the lexical structure of memory. Insofar as words are interwoven with other sorts of representation in memory, such work should be relevant to a very general theory of semantic categorization which encompasses diverse modes of representation.

This paper presents a theoretical framework within which to view the representation of concepts in semantic memory. I believe that it represents a meaningful departure from existing formulations. First of all, the approach presented makes an effort to relate directly the concept acquisition process to ultimate representation of concepts in memory. In fact, the approach was motivated in large measure by concept acquisition research done in conjunction with students in my laboratory.

FL011830

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

John A. Robinson

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

U S DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

Experiments which focused upon acquisition of concepts within certain contexts suggested the importance of such context in ultimate representation of the concept acquired. The investigation of responses to exemplars of a well learned concept indicated that such responses were closely related to the features of stimulus events associated with learning the concept. Such a historical approach opens the way for a detailed study of variables that influence concept representation and can be helpful in accounting for variability in the representation of abstract ideas.

Secondly, the formulation focuses upon the salient exemplars of a concept as a basic unit of conceptual structure. A featural analysis of the concept results from a representation of exemplars as sets of features, be these features "defining" or "characteristic." Thus, newly encountered stimuli may be compared to such structures for purposes of categorization. As I will attempt to show, such an approach may be advantageous in representing ill defined or "fuzzy" concepts as well as those in a continuing state of development. The proposed structure could incorporate diverse features thereby allowing for mixed representation of verbal and nonverbal elements.

Before proceeding to a development of the theory, it might be well to review the major characteristicsof existing theories of semantic categorization as a basis for comparison.

A. Present Theories of Conceptual Representation

A pioneering effort was that of Collins and Quillian (1969). They developed a model called the Teachable Language Comprehender (TLC) to account for processing of linguistic information. The basic building block of TLC is the labeled association. These associations constitute a

network of relationships between nodes of the system. The nodes represent items of information encoded as either a "unit," i.e., the memory's concept of some object, event, idea, etc., or as a "property," i.e., some sort of predication such as may be stated by a verb phrase, relative clause, etc. Properties become linked to concepts and concepts are interrelated in an associationistic fashion. In this way, knowledge is configural with nodes defined only in terms of other nodes in the network. Properties are stored at the highest level possible, a procedure that should result in great cognitive economy. Thus, for instance, the property of having wings would be stored at the node, "bird," rather than at any nodes representing members of that category. In this way, the property is stored at only one place rather than many, and exemplars would have few characteristics directly associated with them in storage.

The theory of Rumelhart, Lindsay, and Norman (1972) is another network theory realized in a program, EL1NOR, in which nodes are connected by specific relationships. For example, the two nodes, "robin" and "bird" are connected by an "is a" relationship (that is, a subset relationship). Here also, semantic memory is conceived of as a network of associations.

Other models have emphasized a set-theoretic representation of semantic memory. Meyer (1970) proposed a model in which individuals are assumed to engage in a serial self-terminating set of comparisons of the name of a subject (S) to names of categories that intersect a predicate (P) as the first stage of a process to verify or disconfirm the assertion that "all S are P." A match verifies that S and P intersect, but a failure to match shows S and P to be disjoint and hence disconfirms the statement, "all S are P." A first stage match leads to Stage 2 in which

comparisons may be made between attributes of S exemplars and attributes of P to see if S exemplars possess all attributes of P. (Other processing mechanisms for Stage 2 were also considered.) Category-size effects were successfully predicted involving increases of number of comparisons in Stage 2. Meyer concluded that stored information about semantic categories included names of categories that they intersect and representations of their attributes.

A related theory is that of Landauer and Meyer (1972). They postulated that when individuals are asked to determine if a stimulus is a member of a specific category, they merely search the set of category members for a match. As evidence for this model, the authors found that reaction time was directly related to category size independent of overall semantic similarity of stimulus to category label.

Smith, Shoben, and Rips (1974) presented a theory with set-theoretic features to account for evidence concerning category structure. In the theory it was assumed that the meaning of the lexical terms could be represented by semantic features. Some of the features were considered essential or defining aspects of the word's meaning (defining features), while others were considered more "accidental" or characteristic aspects (characteristic features). The speed with which a word is classified as a member or nonmember of a specified category was hypothesized to be a function of the "typicality" of the word in terms of its features, both defining and characteristic. The retrieval mechanism postulated possesses two stages. If the word is very characteristic or very uncharacteristic of the category in terms of its total set of features, the decision is made very rapidly. If, however, the word lies between two such extreme

criterion points of typicality, the process proceeds to a second stage in which classification is made with respect to possession of defining features. Thus, the two stage processing model results in the prediction that typicality or representativeness of a word belonging to a given category is inversely related to processing time. The authors characterize the first processing stage as "holistic" in nature, indicating that some overall measure of typicality should be prognostic of classification time.

The authors presented evidence favoring such a processing model. Multidimensional scaling of stimuli for interstimulus similarity and for similarity to a category produced typicality values related to processing time in the predicted manner. Further, a Euclidean metric described the similarity space better than did a non-Euclidean metric. Since there is evidence that Euclidean metrics tend to be associated with holistic or integral dimensional structures (Garner and Felfoldy, 1970), the authors considered the scaling results as supportive of the notion that the first comparison stage is holistic.

The nature of the first processing stage would seem to bear an important relationship to the nature of category structure. Is typicality some additive function of characteristic features, or is it some function of "overall representativeness" of the lexical items? These two formulations and possible variants have divergent implications for the structure of semantic categories.

It is important to note that any nonmember of a category that exceeds the upper criterion point of typicality for that category would "fool" the processing system proposed by Smith, et al. That is, such an item would result in a first stage positive classification. Such

anomalies could exist if the total featural set is not weighted heavily with essential features of the category. Hence, the theory implicitly requires that the internal representation of an item does emphasize the defining features of the category to be matched at any particular time, if it is assumed that classification is usually correct.

In addition to the typicality effects Smith, et al., found a category size effect that is relevant to the assessment of the relative merits of different theories. These authors found that increasing category size decreased "true" reaction times when similarity of exemplars to category was increased by increasing category size. However, increasing category size increased these reaction times when exemplar-category similarity was decreased by increasing category size.

For example, the reaction time of "butterfly" for smaller category, "insect," was shorter than for larger category, "animal," where "butterfly" was more similar to "insect" than "animal." In contrast, reaction time of "cantaloupe" for smaller category, "melon," was larger than for larger category, "fruit," where category size-similarity relationships are reversed. Such a result seems to be in keeping with the typicality aspects of the Smith, et al. theory. In contrast, a network theory, such as that of Collins and Quillian, postulates that larger categories in the hierarchy are represented by nodes farther in the network from that of an exemplar than nodes for smaller categories. Thus, increasing category size should result in lengthening the path to be searched and hence a longer reaction time in all cases. The results are thus considered a blow to this type of network theory. Meyer's (1970) theory postulates a priority of search of categories close to the size of the probe in the first stage,

hence, his theory appears to have the same problem as that of Collins and Quillian in this respect.

Insofar as the network theory leads to a more "associative" interpretation of semantic structure, the Smith, et al. data would cast doubt upon such a viewpoint. However, it is possible that interpretations of exemplar-category relationships in terms of straightforward associative links is not precluded.

Other investigators have presented evidence bearing upon the relationship of members to categories.

Rosch (1975) used superordinate semantic category names to prime comparison of exemplar pairs for the category. She found that priming facilitated "same" responses to physically identical good members of a category and hindered these responses to physically identical poor members of the category. Since line drawings were used to represent exemplars, this "goodness of exemplar" effect applies to nonverbal stimuli. Subsequent experiments were performed, using verbal stimuli also, to show that the same kind of information processing occurs at the level of meaning. The results seem to support a "family resemblance" interpretation of category membership as advanced by Wittgenstein (1953). Wittgenstein and other philosophers have noted that few natural language categories have well-defined sets of inclusion criteria. Such categories may be viewed as "fuzzy sets." Indeed, a fuzzy set theory approach has been applied to deal with such vagueness in natural language (e.g., Hersh and Caramazza, 1976). I shall return to this approach shortly.

Loftus (1975) has countered the family resemblance interpretation of Rosch with one based on the notion of spreading activation. According to

such a view, a prime sends its activation over a network of associations in which stimuli would receive more or less activation according to associative strength within the network. Thus, good members would receive more activation than poor members by virtue of perhaps more and closer links with the prime. Good members would then be encoded faster than poor members and produce the faster response time observed.

A clear differentiation between the two types of postulated structures may not be possible in the type of experimental situation presented.

It is clear that the data on categorization defy the Aristotilean notion of clear category boundaries. Lexical terms do not seem to be associated with enough distinguishing features in natural language usage to enable one to make category inclusion decisions with such great precision. This fuzzy aspect of natural language categoriesseems to contrast sharply with the clear distinctions derived from artificial classifications of concept learning. What is the nature of the fuzziness of natural language categories? Wittgenstein (1953) pointed out that an examination of a category structure associated with a given term will reveal "a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities, sometimes similarities of detail (p. 32e)." The word, 'game," is used as an example. There are Olympic games, tennis, childhood games, etc. One may have a class of games in which a clearly competitive aspect and rules serve as defining features. However, the incidental playful aspects may be focused upon to define another class of games. Thus, certain characteristic but non-defining features embodied in one well-defined class may be used to define another subcategory associated with the same label. In this way no clear

defining characteristic is true of all members of the labeled category. Olver and Hornsby (1967) have in fact dealt with such "complexive structures" in the categorization behavior of young children. They identified "rope match" and "key ring" groupings as examples of the processes by which family resemblance concepts of the sort described above could develop. Such concepts could in fact be viewed as disjunctive concepts, and, in this sense, fall within the scope of concept learning paradigms.

The notion that word meanings, which in turn define category membership, are generally not precise can be contrasted with the view that such meanings can ultimately be represented by some finite set of semantic features. Katz and Fodor (1961), for example, developed a theoretical point of view in which disambiguation of word meanings within given contexts takes place through "markers" which represent implicit features associated with a word in its use. These markers enable one to discriminate among different senses of the word. Difficulties arise, however, in attempting to specify exhaustively these distinguishing features. Bolinger (1965) has illustrated how an attempt to ascertain the features necessary and sufficient to find the sense of a word in its various uses become a virtually impossible task. For example, the word, "bachelor," may refer to an unmarried male, a young knight in certain circumstances, a young fur seal, or someone receiving a bachelor's degree. A given context may contain implicitly certain features which eliminate one sense or another. For instance, "a bachelor's speaking," eliminates the sense associated with fur seal, because "capacity to speak" would be an inferred distinguishing semantic feature. If all contexts are

considered, it would be necessary and sufficient to specify completely the meaning of a word in terms of the instances which it includes.

The above considerations suggest ways in which category structure could be imprecise. Linguistic category structures no doubt develop in large measure through long-term experience within a given context. Degree of precision in assignment to categories may well be dictated by a user's needs within a context. For example, the term, "bachelor," in its more usual sense is generally applied to individuals within a given range of descriptive values. Within this domain the term may have precise meaning, i.e., "an unmarried male." If a man has never been married, and he is within a certain age range in which specification of his marital status may be a topic of some concern, classification of an individual meeting such specifications would become precise. However, classification behavior may become less precise when attribute values lie outside the given domain. Unmarried men outside of a certain age range who have been married may be classified with less certainty. (Men who have been previously married are often classified in other categories. Thus, there is a lowered probability of association with the label "bachelor.") Categories may then be quite precise within an appropriate experiential domain but imprecise otherwise. Since such experiential domains are often in a state of flux, individuals must have mechanisms for incorporating novel inclusion relationships into their systems of representations. A theory of conceptual representation would seem to require a means for storing information that is potentially relevant to deciding about class inclusion relationships, because individuals are capable of making inferences about exemplars outside a usual context.

The decision process would then require flexibility in generating a representation appropriate to inferences in new contexts.

Such information could be considered "irrelevant" to decisions made within the meaning criteria adequate for prior decisions. In the above example, classification of formerly married men with respect to bachelorhood would require some representation of formerly unused characteristics along with some related decision procedure. An adequate theory should account for a representation process flexible enough to accommodate the requirements of continuing inference and revision of concepts.

Another area of research which bears upon understanding of category structure involves abstraction of prototypes in pattern classification tasks. Posner (1969) developed a prototype distance model that predicts that when a series of similar visual stimuli is experienced, a prototype is abstracted at the point in the multidimensional similarity structure which represents a central point of similarity to all stimuli. He and his associates have presented evidence to support this point of view (e.g., Posner & Keele, 1968). In contrast, the attribute-frequency model (Neumann, 1974) predicts that a prototype is formed as a pattern composed of the most frequently experienced elements on each dimension of variability. Evidence has been presented that is supportive of this point of view (e.g., Barresi, Robbins, & Shain, 1975). Neumann (1977) in a series of experiments produced no conclusive evidence with regard to alternative points of view, but pointed out further complexities including measurement problems. This area of research has involved a stimulus domain that may yield results not representative of a broad range

of symbolic processes. For example, most representations of categories may encompass complex interconnections of verbal and imaginal components that may not be associated with stimulus sets of random dot patterns. It is also likely that the notion of a perceptual prototype of the sort proposed by Posner will be difficult to apply to complex representations that include dimensions not comparable in units of measurement, or, indeed, not quantified. Nonetheless, categorization of perceptual patterns must be accounted for in a complete theory of conceptual representation.

The work cited above has dealt largely with uncertainties in category boundaries. As mentioned earlier, a fuzzy set theoretical approach has been applied to the treatment of vagueness of natural language concepts. Theorists such as Labov (1973) and Lakoff (1973) have proposed that a fuzzy logic must apply in dealing with meaning criteria in natural language. Investigators such as Hersh and Caramazza (1976) and Zadeh (1965, 1973) have developed quantitative approaches relevant to imprecision in class inclusion relationships. In conventional set theory, a membership (characteristic) function specifies which elements are members of a set, i.e., $x \in X$ if $f_X(x) = 1$, and $x \notin X$ if $f_X(x) = 0$. Thus, the usual dichotomy in truth values applies. In fuzzy systems, the degree of membership and corresponding truth value of the proposition, $x \in X$, may be any value in the closed interval [0 to 1.0]. This quantification of the function is then used to define operations on fuzzy sets analogous to those in conventional systems.

In the context of linguistic relationships, language operators such as adverbs and adjectives, are considered as operators on fuzzy sets through the effects imposed upon characteristic functions associated with

terms modified by such operations. The details of the formal system need not be considered here. However, Hersh and Caramazza (1976) have shown that responses of subjects to inclusion relationships can be interpreted in terms of such fuzzy logical operations.

The possibility that a fuzzy logical system is descriptive of conceptual representation is important. However, it is necessary to explore the mechanisms intrinsic to such representations. What are the factors that determine the value of a characteristic function? How are these functions modified in relationship to various inclusion relationships? It will be necessary to relate quantitative relationships of the descriptive theory to mechanisms underlying such relationships. First, evidence from studies of concept attainment will be considered, with emphasis upon acquisition factors that might relate to ultimate representation of a concept.

B. Evidence from Concept Attainment

The concept attainment paradigm has not been used extensively to assess variables important in conceptual structure, probably because hypothesis theory, which motivated a great deal of such research, was not primarily concerned with the topic. Hypothesis theory (e.g., Bower & Trabasso, 1964; Falmagne, 1970; Gregg & Simon, 1967; Levine, 1966; Nahinsky, 1970; Restle, 1962; and Wickens & Millward, 1971) has taken an "information filter" approach, generally assuming that individuals abstract dimensional information from stimuli while disregarding the specific stimuli during the process. Katz and Denny (1977) recently showed that concrete concepts are more readily attained than abstract concepts, largely because previous concrete instances are more readily available from memory than abstract instances. This work, which controlled for other

variables such as frequency and meaningfulness, corroborates that of a number of others (e.g., Baum, 1954; Dattman & Israel, 1951; and Heidbreder, 1946). Such results suggest that learning context is an important variable in determining the course of learning. Thus, it is possible that such context influences ultimate representation of the concept.

Mahlosky, Oeschger, and O'Leary (1977) reported concept learning experiments that placed primary emphasis on the influence of context upon concept attainment. Stimuli were used that consisted of descriptions of individuals on four dimensions. If the subject is presented with the descriptive characteristics alone, as was true for one condition of the first experiment, he may associate an image or some descriptive representation with the characteristics to produce a unitized and distinct picture of the stimulus event. In another condition of the first experiment the verbal description was associated with a known exemplar, e.g., Custer--a real, American, historical soldier. In this instance context and distinctiveness were maximized, with context known to the subject to be irrelevant to task solution increasing the uniqueness of individual stimuli. In contrast, stylized stimuli consisting of combinations of abstract symbols were used for another experimental condition. In all stimulus conditions problems were varied with respect to the minimum number of stimuli necessary to identify solution. An examination of correct response probabilities prior to the last error revealed that context enhanced memory for specific category classifications. However, context adversely affected efficiency of concept attainment, but only in conditions in which many stimuli were necessary to identify solution. The finding thus indicates that highly contextualized stimuli encourage a strategy in which

individual stimulus classifications are stored and contrasted in memory during learning. When many such classifications are necessary to specify the concept, the memory load required for such a strategy impedes learning.

It is important to note that results held for the pictorial stimuli despite the abstract nature of the dimensions, whose values had to be accessed from an information store separate from that derived from a visual representation that exemplifies these values. If one sees a representation of Galileo, information about era and occupation is separate from information about features such as beardedness. Thus, exemplars give no direct task relevant information. It is reasonable to hypothesize that such contextual influences in the history of concept acquisition could make themselves felt in the representation of the concept in memory. Distinctive context of essentially mnemonic value could lead to the representation of the concept through very distinct examples, regardless of typicality of component dimensions.

To investigate this possibility, Nahinsky, Oeschger, and Morgan (unpublished experiments) have done work to assess the influences of learning context upon classification performance with respect to a well learned concept. Individuals were required to learn simple one-cue concepts with the highly distinctive person stimuli described above. They were presented with three cycles of eight feedback person descriptions interspersed with four no-feedback trials to check on learning by memorization. In one condition each random permutation of the stimuli had a change of person-exemplar for each stimulus. For example, if the person description, "real, American, historical soldier," was exemplified by Custer on one cycle, the same description might be exemplified by Sherman on the next

cycle. In another condition the exemplars were not changed from cycle to cycle. The subjects were instructed that only the descriptive dimensions were relevant for solution, and hence the two conditions were informationally equivalent from the standpoint of solution. Criterion responding was generally well established by the end of the three cycles, at which point a test series of 16 trials was introduced with no feedback. Eight of the stimuli were "old" stimuli, i.e., stimuli encountered during learning, while the other eight were "new" stimuli. New stimuli were matched with old stimuli on descriptive combinations. Thus, the new set differed from the old set only with respect to the contextual exemplars.

It was found that in the condition involving no change in associated exemplars over trials, the response latencies were significantly shorter for old stimuli than for new stimuli. All results held when corrections were made for problems not solved. Therefore, it appears that distinctive associative context was a prime factor in facilitating the storage and retrieval of correct classifications. This relationship indicates that accessing and classification are not in terms of defining attributes alone, in which case old and new stimuli should not have been distinguished. In addition, distinctive associative context apparently played a large role in developing conceptual representation.

In a second experiment the constant context condition of the above experiment was used. However, new stimuli varied with respect to the frequency with which irrelevant values appeared on positive instances during learning trials. Each new stimulus was scored for irrelevant dimensional overlap with positive instances on learning trials. Thus, if three of the learning positive instances were soldiers, a new stimulus

with value, "soldier," would get a score of three on the occupational dimension. Adding scores for all three irrelevant dimensions produced the irrelevant dimensional overlap score. New positive instances with irrelevant values more frequently presented during learning were classified significantly more rapidly than new positive instances with less frequently presented irrelevant values. However, both of these sets of new stimuli were classified significantly less rapidly than old positive instances. The results of the two experiments thus suggest that concepts are represented in memory through highly associated acquisition context of a holistic nature and to a lesser degree through isolated characteristic features associated with acquisition.

An analysis of the high overlap new PIs in the second experiment revealed that each descriptive combination was represented on one old PI. Thus, if an old PI was Sherman, a new PI in the high overlap category might be Custer, another real, American, historical soldier. Such a correspondence suggests that the high overlap superiority for new PIs may have been mediated by some high association of a new PI with a highly related old PI. To test this possibility, the experiment was repeated with stimuli including only the verbal descriptive values. Naturally, no new PI could be the same as an old PI on the set of descriptive dimensions without being identical to it in this case. Thus, overlap relationships differed, though very slightly, from those of the prior experiment. In this experiment, the high overlap low overlap difference disappeared for new PIs. However, the old PIs were responded to significantly faster than were all new PIs. Thus, it would appear that the effect of the overlap variable in the previous experiment was mediated by associations between

highly similar stimuli. The old-new difference coupled with the divergent results of the two experiments for the overlap variable suggests that the superiority of old stimuli in the present experiment results from the unitization of old descriptive combinations into some personification, real or hypothetical, which can be associated with the category. Such a hypothesis is consistent with prior evidence (Mahinsky, et al., 1977). The evidence of these experiments is in accord with the model to be proposed. An initial search of the set of representative exemplars appears to be involved.

C. Well Established Versus "Open" Conceptual Systems

It can be seen that a good deal of present theoretical work is designed to account for categorization according to at least reasonably well defined and established "closed" concepts. Thus, the ability to categorize "chair" as an item of "furniture" is conceived of as a process whereby features are matched to a category definition, or well established associations are called into play. A historical approach, implied by use of the concept attainment paradigm as presented above, can examine mechanisms whereby new information is incorporated into a structure and new examples of a concept are dealt with. Assessment according to static categories leaves a wide range of conceptual classification unaccounted for. Many concepts remain in a state of development and refinement over a long period of time, i.e., they are "open." Their application is dependent upon the context within which they must be used during such a period. As Kintsch (1974) has pointed out, ". . . Memory must have the capacity to make inferences. The fact that one can explain this or that experimental phenomenon without recourse to inferential process merely

means that one deals with data sets that are impoverished when taken out of context (p. 202)." A given lexical item may be referred to a number of contexts, each associated with its categorical structure and developing in its own manner. Thus, "gentleman" may be referred to a class of individuals possessing certain central personal qualities or to a class of individuals distinguished by certain conventional behaviors. Each class, of course, has its own structure, and, once the subtle shades of context have been determined, processes of retrieval of semantic information instituted by the word, "gentleman," will follow the course dictated by the context. The approach to be presented attempts to deal with such contextually determined structures.

B. Theoretical Framework

Any theory of conceptual representation should account for established findings as well as for its specific predictions. Post hoc handling of results is, of course, not a very impressive accomplishment for a theory, but certainly represents a minimal starting point. The formulation to be presented will take account of results in the realm of statement verification as well as results cited above for concept attainment. The theoretical framework will be developed in two stages. First of all, a scheme for conceptual representation will be elaborated. Secondly, procedures for comparing stimuli with such representations and retrieval of information from them will be considered. Such procedures are no doubt determined by contextual demands and will vary with the tasks given to an individual. Hence, it may be well to consider such procedures as variable means for studying the underlying conceptual structure.

As it has been suggested, it is proposed that concepts, at least

during their developmental stages, can be represented through a series of closely associated exemplars. These exemplars, in turn, can be represented as sets of features. Thus, the representation takes the form of intersecting sets of features. Insofar as such a representation can be used for a featural analysis, predictions based on this formulation may be in accord with a number of predictions based upon current featural theories of statement verification. The proposed framework also resembles theories such as that of Landauer and Meyer (1972) in that categories are represented by some set of exemplars. However, it should be possible to derive appropriate differential predictions, which will be the task to be dealt with in Section E. At any rate, the concept attainment work cited gives some credence to the approach presented here.

1. Representation system. Assume that a concept and associated concept (or category) name is represented by some highly associated set of K exemplars. That is, the name associated with the concept results in independent retrieval of each of the K exemplars with some high probability, p. Thus, if K′ represents the subset of exemplars sampled, it is binomially distributed with

$$E(K') = Kp$$
$$\text{and } \sigma^2(K') = Kp(1-p). \qquad (1)$$

Each retrieved exemplar, i, results in the construction of a set E_i of features relevant to the domain in which the concept is to be applied. That is, the featural representation of the exemplar, as well as the exemplar, is associated with the concept name. Consider

$$\bigcap_{i=1}^{K} E_i = H, \qquad (2)$$

a population value, which will be called the set of defining features. Then let,

$$n(\mathfrak{n}) = n. \qquad (3)$$

In this context n becomes the number of defining features. As it can be seen, a defining feature is considered to be a feature common to all exemplars in the set. However, in instances in which concepts are not completely developed, the status of a feature with respect to "definingness" may vary over time. Further, let K' be the number of exemplars retrieved for an application of a concept, where, of course, $n(K') \leq n(K)$. Then,

$$\bigcap_{i=1}^{K'} F_i = \mathfrak{n}', \qquad (4)$$

$$n(\mathfrak{n}') = n'$$

are the functional set of defining features and the functional number of defining features respectively for the application.

In the above context $N = n(\bigcup_{i=1}^{K} F_i)$ is the total number of characteristics which might be considered, and N' is the corresponding functional number with respect to given K'. Consider some subclass, A_m, of K-1 exemplars chosen from the K exemplars, and order the featural representations for this subclass, such that the jth representation of the subclass is denoted $F_{j,K-1,m}$. Then,

$$\bigcap_{j=1}^{K-1} F_{j,K-1,m} = D_{K-1,m}. \qquad (5)$$

Next, consider all $\binom{K}{K-1}$ subclasses, A_m, that can be formed from the K exemplars. Then,

$$\bigcup_{m=1}^{\binom{K}{K-1}} D_{K-1,m} = A_{K-1} \qquad (6)$$

is the set of all features common to at least K-1 of the exemplars. It follows that

$$n(A_{K-1}) - n = n_{K-1} \qquad (7)$$

is the number of features common to exactly K-1 exemplars. Define $A_{K-2}, A_{K-3}, \cdots, I$ as the number of features common to K-2, K-3, ---, 1 exemplars respectively. Then, following procedures and notation in a manner analogous to that of (5), (6), and (7),

$$\bigcup_{m=1}^{\binom{K}{K-2}} D_{K-2,m} = A_{K-2},$$

$$\bigcup_{m=1}^{\binom{K}{K-3}} D_{K-3,m} = A_{K-3}, \qquad (8)$$

$$\vdots$$

Where,

$$D_{K-2,m} = \bigcup_{j=1}^{K-2} E_{j,K-2,m},$$

$$D_{K-3,m} = \bigcup_{j=1}^{K-3} E_{j,K-3,m}, \cdots.$$

Noting the transitive subset relationship, $A_{K-1} \subset A_{K-2} \subset A_{K-3} \subset \cdots \subset A_1$, it can be seen that

$$
\begin{aligned}
n(\Lambda_{K-2}) - n(\Lambda_{K-1}) &= n_{K-2}, \\
n(\Lambda_{K-3}) - n(\Lambda_{K-2}) &= n_{K-3}, \\
n(\Lambda_{K-4}) - n(\Lambda_{K-3}) &= n_{K-4}, \\
&\vdots \\
&\vdots \\
&\vdots \quad ,
\end{aligned}
\tag{9}
$$

where n_{K-1} is the number of features common to exactly K-1 exemplars. We may then allow n'_{K-1}, n'_{K-2}, n'_{K-3}, - - - to denote functional values corresponding to values shown in (7) and (9).

By (7) and (9), and noting that N - n is the number of defining features,

$$
p_{K-j} = \frac{n_{K-j}}{N-n} \tag{10}
$$

is the probability that exactly K-j exemplars share an attribute that is nondefining, with $p_{K'-j}$ a corresponding functional value. Further, let the variable K-j be denoted by X with corresponding functional value, K'-j denoted by X'. Then,

$$
E(X) = \sum_{j=1}^{K-1} (K-j) n_{\frac{K-j}{N-n}}, \tag{11}
$$

the expected number of exemplars per nondefining feature. If equalities in (7) and (9) are substituted,

$$
E(X) = \sum_{j=2}^{K-1} (K-j) \frac{[n(\Lambda_{K-j}) - n(\Lambda_{K-j+1})]}{N-n} \tag{12}
$$

$$
+ \frac{n(\Lambda_{K-1}) - n}{N-n} = \sum_{j=1}^{K-1} \frac{n(\Lambda_{K-j})}{N-n} - \frac{n}{N-n} .
$$

.Thus, E(X) can be expressed as a linear function of the sum of the $n(\Lambda_{K-j})$, the number of features shared by at least K-j exemplars. The value of E(X′) can be found as follows:

$$E(X') = \sum_{K'=0}^{K} E(X'|K')\, p(K')$$

$$= \sum_{K'=0}^{K} \left[\sum_{j=1}^{K'-1} \frac{n'(\Lambda_{K'-j})}{N' - n'} - \frac{n'}{N'-n'} \right] p(K'), \qquad (13)$$

where $p(K') = \binom{K}{K'} p^{K'}(1-p)^{K-K'}$

It would be helpful to make some simplifying assumptions about the distribution of X. Let g be the probability that any of the X exemplars possess some feature. This value can be estimated from data on observed probability of featural possession. If we can assume that possession of the feature by one exemplar is independent of its possession by another, then the probability that exactly K-j exemplars show a nondefining attribute is

$$\binom{K}{K-j} g^{K-j}(1-g)^{j} / (1-g^{K}), \quad 0 \leq g < 1,\ j=1, 2, \ldots K-1 \qquad (14)$$

according to the binomial distribution and noting that the probability a feature is nondefining is $(1-g^K)$. From (14) it follows that

$$E(X) = \frac{Kg}{1-g^{K}} \left(\frac{1-g^{K-1}}{1-g} \right)$$

$$\sigma^2(X) = \frac{Kg}{(1-g^{K})} \left\{ (K-1)g(1-g^{K-2}) + (1-g^{K-1}) \left[1 - \frac{Kg\,(1-g^{K-1})}{(1-g^{K})} \right] \right\}, \qquad (15)$$

To get comparable values for X', we note that

$$E(X') = \sum_{K'=0}^{K} \sum_{x'=0}^{K'-1} X' \, p(X'|K') \, p(K'), \qquad (16)$$

and substitute appropriate values for X' and $p(X'|K')$ from (14) while noting the binomial distribution of K'. We get

$$E(X') = \sum_{K'=0}^{K} \frac{K'g(1-g^{K'-1})}{(1-g^{K'})} \binom{K}{K'} p^{K'}(1-p)^{K-K'}. \qquad (17)$$

The equation, unfortunately, does not permit further simplification. However, for reasonably small K, computation should not be too difficult. Further, $E(X)$ may be a reasonable approximation for use if an individual retrieves a small number of high strength exemplars.

The simplifying assumptions may, of course, not do justice to the true state of affairs. However, for present purposes, they should suffice to illustrate the approach. At any rate, assumptions could be tested empirically, and more adequate distributions could be substituted.

As noted above, there are N-n nondefining attributes. Then,

$$A = (N-n) \; E(X) \qquad (18)$$

becomes a measure of total occurrence for nondefining attributes. Also,

$$B = nK \qquad (19)$$

is a measure of total occurrence for defining attributes. Corresponding expectations of functional values could, of course, be substituted. However, (18) and (19) may be assumed adequate for present purposes.

For a poorly defined concept, A should be large relative to B, reflecting a predominance of nondefining or "incidental" features in the total representation. Conversely, if B reaches a large value relative .

to A, the concept approaches a clearer definition, with greater emphasis on common features in representation. A potential anomaly should be noted in the situation in which g approaches unity. At that point all features become common. Therefore, the restriction, $0 \leq g < 1$, in (14) is required.

Let α and β be appropriate monotonic functions of A and B respectively. Then, the above considerations can be reflected in an "index of fuzziness,"

$$F = \frac{\alpha(A)}{\alpha(A) + \beta(B)} \quad . \qquad (20)$$

The functions, α and β, may be linear, exponential, or any of a number of monotonic functions that must be determined by empirical fits of the data. However, the general relationships involving A, B and conceptual ambiguity must hold for the scheme to work.

What are the ecological advantages of such a representation structure for the individual? After all, a great deal of "irrelevant" information must be stored. A representation system that relies exclusively on featural analysis would appear to have great advantage in terms of storage and retrieval.

However, such a completely logical featural analysis does not take into account the tentative nature of human knowledge. The individual cannot be certain about the nature of a concept that is evolving inductively in his experience. It is possible, for example, that some encountered exemplar of a concept may contain no features common to all previous exemplars. A "successful salesman" may possess no qualities shown by all previous "successful salesmen." Thinking would have to be

revised, and features previously thought to be incidental to the concept would have to be reconsidered. The concept may prove to be disjunctive, or it may resemble complexive structures referred to earlier. Such possibilities could lead to a split of the representation into component representations. A system that enables one to use a limited number of exemplars to redintegrate a large amount of featural information may be quite efficient. Different representations of the same exemplar within different contexts may also prove to have advantages. In addition exemplars common to a number of concepts may link them associatively.

Interesting possibilities exist for conceptual development. With increasing experience, the representation may be truncated. Certain features that are shown to be nondiscriminating with respect to the application contexts may be eliminated from exemplar representations. Thus, commonality would increase so that, in the limiting case, all features would be common. At that point all exemplars would match each other in representation, and, perhaps, only one would suffice as a retrieval cue. Such a result might be considered as a manifestation of a "consolidation" process in which elements of a structure become integrated into a single cognitive unit of representation. Hayes-Roth (1977) has conceptualized the development of such units, which she labels "cogits." Such units may be analogs of Hebb's (1949) cell assemblies.

Paradoxically, at some advanced state of development, reduction of cues would produce high weights in a number of "nondefining" cues. Reduction of cues of lesser commonality would increase g for those cues remaining. Thus, an increase in ambiguity of the concept would result before a final resolution. This increased uncertainty should manifest

itself in an individual's showing more "thoughtfulness" in application of the concept with increasing knowledge at some point in its development. Perhaps, people less naive about a concept are less prone to make "snap" judgments about its application. At any rate, this possibility is testable.

The proposed formulation should be contrasted with the notion of a prototype. Rather than postulation of some "idealized" exemplar of a concept, multiple representations linked by means of contextually relevant features is proposed. There are certain advantages to the latter approach. For one thing, the necessity of somehow averaging dimensions of representation to arrive at some modal representation is removed. Such an averaging process could prove difficult when disparate dimensions are incommensurable or, indeed, there are dimensions that do not provide a ready metric. It is also the case that the formulation proposed in this paper does not require one to confront the problem of representation of an abstract idea by some single concrete image, a notion philosophically unsatisfying to many.

2. Stimulus - concept comparison processes. Processes involved in applying a concept to specific events are no doubt varied and are dependent on many aspects of the situation in which the concept is to be applied. There will be no attempt to catalog completely the possible comparison processes. However, I will attempt to deal with major aspects of processes which involve comparison of a stimulus with a conceptual category structure in order to determine if it is a member of the category. The major concern will, of course, be the extent to which such comparison processes can reveal the underlying conceptual representation

process. Further, manipulation of variables associated with such processes should provide a means of discriminating among postulated conceptual representations.

First, let us assume that when a stimulus is presented, and an individual is asked to verify its membership status with respect to some conceptual category, he constructs some representation of that stimulus to be compared with the conceptual representation. The representation can take the form of an image, a network of verbal and symbolic descriptions, or some combination of these two modes of representation. Rosenberg and Simon (1977) have, in fact, presented a model which does allow for integration of elements for such diverse modes of representation. Regardless of the mode of representation, it seems reasonable that such an entity is translated into some feature set for comparison with the conceptual representation. Features may well be correlated in a "holistic" representation as suggested by results cited above for Smith, et al., (1974). In any case, the nature of featural comparison is of central concern.

Certain dimensions along which comparison processes can vary will be examined. First of all, it may be well to consider if the processes involve stages, starting with some "global" process for clear cases and extending to a more analytic process if finer discrimination is required. The Smith, et al. model provides for two stages, with second stage comparisons on a featural matching basis as has been noted. Within the present context the first stage might take the form of a search of the exemplar set of the concept representation to see if the given stimulus is itself in the set.

Thus, a highly associated member of the category would be classified immediately. It is also possible that such a first phase would incorporate a search of the exemplar set for one member highly similar to the stimulus to be compared. In this case, a very representative example of the category could serve as the basis of featural comparisons similar to those postulated by Smith, et al. As it will be seen in Section F, a first phase search to see if the stimulus is a member of the representation set has much to recommend it.

Next, consider the possible mechanisms for featural comparisons per se. There may or may not be an initial holistic comparison phase. However, at some point a feature matching process would probably have to be instituted for a certain proportion of stimuli even if a holistic stage is allowed. (The Smith, et al. model, of course, calls for a feature-matching process from the start in all cases.) If feature matching is called for, how does the process take place? Is it based solely on the total number of features represented, or is it based also upon the degree of representativeness of features as determined by a measure such as F shown in (20)? The latter possibility is more likely for a representation based upon a number of concept exemplars in which feature frequency can be assessed. In any case, it is necessary to consider some feature list derived from the stimulus representation with respect to appropriate dimensions. Label this domain the comparison domain. Suppose E_s is the set of features comprising the representation of stimulus S. Then the comparison domain is $E_s \cap (\cap \bar{F}_0)$, which we denote as E. Here $\cap$ is as defined in (2), and F_0 is the set of features with respect to the array of dimensions in E_s possessed by no stimuli in the

representation set. Thus, $F_s \cap F_0$ is the set of features possessed by the stimulus, but not by any stimulus in the representation set. Within E we can derive formulas analogous to those represented in (2) through (20). For example, N_s is the set of all features common to the representation set as well as the comparison stimulus, and F_s is the index of fuzziness relative to comparison with S, if E_0 is ignored. The role of E_0 presents a problem for considering comparison strategies. The set of features not found in any member of the exemplar set should serve as a "counter-weight" to N_s. In this respect, ambiguity of the comparison might be reduced by E_0. However, both N_s and $n(E_0)$ may be large resulting in conflicting tendencies in the decision process. For the moment, suppose E_0 is small. Then assume it is weighted as a set of features possessed by one exemplar in determining F_s for the comparison. In this case E_0 tends to make a match to the category more difficult, a state of affairs which seems reasonable for relatively small E_0. The "large E_0" case will be considered subsequently.

It is apparent that many variations in feature processing are possible. Suppose the process involves a comparison with respect to defining features and that matches depend upon finding correspondences between stimulus and concept for such features. If N_s is the total number of features common to exemplar and comparison stimulus with n_s the number of "defining" features, ease of classification should be roughly proportioned to $\frac{n_s}{(N_s + E_0)}$, the probability of a defining feature match if any one feature is considered. For a looser criterion of matching,

the ratio might be $\frac{n_s + (N_s - n_s)\gamma}{(N_s + F_0)}$, where $N_s - n_s$ is the number of "characteristic features" weighted by some constant γ, $0 < \gamma \leq 1$, to account for the use of such features in making a match. Such a formulation would be closer to that of Smith, et al.

The relationship shown above would not account for possibilities within the representation structure proposed in this paper. Sets of exemplars of different concepts may produce like values of N and n as well as N_s and n_s for given comparisons. However, these sets could have quite different structures and correspondingly different values of F_s for such comparisons. A highly variable set with respect to nondefining features would produce a value of F_s quite different from that produced by a set of exemplars all quite similar with respect to nondefining features, even though both sets possess identical N_s and n_s values for given comparisons. It is possible that the F_s would be more prognostic of classification ease than some measure based only on feature counts.

The actual feature matching process can occur in a number of ways. However, detailed specification of the process may not be necessary to derive meaningful tests of theories. Search of features may be serial or parallel, but either type of search can yield predictions that preserve relationships stated in preceding paragraphs. Certainly, the problems of distinguishing empirically between serial and parallel processes are well documented (e.g., Townsend, 1972). Search may be made until some critical number of matches on defining features is made or until a complete test for all defining features is made. Certain combinations of defining and characteristic features in toto may be

required. In general, predictions based upon different matching criteria should be correlated. The use of F_s would incorporate various processes involving a defining and a nondefining feature search. By weighting features for degree to which they are representative within the representation set, total ambiguity is specified in any comparison, and some search of features would generate a measure resembling F_s. Section E will deal with the problems of differential predictions for such comparison processes.

The problem of large F_0, alluded to earlier, must be considered. If a stimulus possesses many features possessed by no members of the representation set with respect to dimensions used in the representation set, dissimilarity of stimulus to concept would be quite great. If F_0 reaches a critical size relative to F, one might expect very rapid classification in the nonmembership category. When the infrequent atypical stimulus occurs, i.e., one that is very much unlike members of the category to which it belongs, there would be a false negative classification.

Another possibility presents a more interesting problem. Suppose the representation of a stimulus to be compared with a concept representation contains features on a number of dimensions not used in the concept representation, but the stimulus is yet an exemplar. That is, the stimulus is represented by elements well outside the domain in which the concept has been encountered. The example cited earlier in which a previously married man is evaluated for his eligibility for the bachelorhood category might serve as an illustration. Such an individual possesses characteristics on dimensions not usually associated with a representation of the concept.

In such a case the comparison process would be weighted heavily with features which would not allow a comparison, a state of affairs that should materially affect the decision process. It might be necessary to assume that if the number of features in the stimulus representation outside the domain of concept representation reaches a critical proportion of the total for the stimulus, an individual must revert to a more detailed and deliberate search for matches on defining features. In any event, the comparison process should be impeded, and this state of affairs should be reflected in data.

The possibilities presented above, although not exhaustive by any means, suggest relevant tests of hypotheses about comparison processes with respect to concepts. In the next section a general outline of possible empirical approaches will be considered.

E. Empirical Tests of the Representation System

In deriving tests of the representation system it is, of course, important to consider situations that would differentiate empirically the present approach from others. The proposed system and other approaches can be expected to generate the same predictions in a number of situations.

This assertion follows from the fact that, in the proposed system, exemplars are used to generate featural representations, and hence exemplar and featural sets should be correlated. However, the system proposed allows for variations in the way featural representations are constructed which can lead to predictions not dealt with by other theories. For example, suppose category set size is held constant while overlap of members within categories is varied. Differential effects of

such variation in overlap could not be accounted for by a model such as that of Landauer and Meyer (1972) which involves only a search of the set of exemplars. Conversely, suppose featural similarity of exemplars to category is held constant, while the structure of the category with respect to overlap relations among salient exemplars is varied. The effect of such structural variations should provide difficulty for a theory based only on feature comparisons.

In order to develop empirical tests of conceptual representation, concepts to be represented must be subjected to scaling along a number of dimensions. It is not necessary to consider approaches to such scaling in detail here because of the volume of work in existence on this general problem. Subject based measures of category size and composition seem appropriate. Cohen, Bousfield, and Whitmarsh (1957) provide perhaps one of the best examples of norms which could enable one to estimate category size and composition. These norms were derived by asking 400 subjects to list four items under each category heading. In this way, one could assess each category in terms of its most salient exemplars. If one uses some threshold frequency criterion, it would be possible to catalog conceptual categories in terms of size and composition of salient exemplar sets. There have been arguments about the "true" measure of category size (see Wilkins, 1971). However, within the present context, it appears that a measure of category size of the sort derived by Cohen, et al. is called for. Rated category size as obtained by Battig and Montague (1969) would not provide the information necessary to derive an empirical test.

Given measures of category size and composition, concepts could be matched for abstractness using analyctic measures (e.g., Landauer and Freedman, 1968; Meyer, 1970; Landauer and Meyer, 1972). If two concepts produce the same set size measure, but one is more abstract than the other in that it contains a number of nested categories, assessment of conceptual structure in terms of features might be confounded with a subcategorization process. Thus, control for abstractness is necessary.

After an assessment of category membership in terms of salient exemplars, these exemplars could be evaluated for characteristic features. Subjects could be asked to list the features of given exemplars with techniques of the sort used by Cohen, et al. for assessing category membership. In this way the category structure of the type proposed here can be assessed for various concepts. The ease of comparison of stimuli with categories can be used as a measure to compare theories.

Numerous possibilities exist. As suggested above, category size could be held constant, while featural overlap is varied. Also, featural set size could be held constant, while category overlap structure is varied. In general, if exemplar overlap proves to have a significant effect upon comparison behavior independent of category size and featural similarity of stimuli to category, the need for a representation process of the type proposed would be indicated. Of course, the structural measures proposed here should be predictive of results. In general, the effect of exemplar overlap should hold over a range of comparison strategies if the proposed system has validity. However, different comparison strategies may be linked to correspondingly different measures as suggested in Section D2.

It is important to note that the domain of concepts used for empirical test can well be extended beyond concepts of the type usually studied. Conceptual categories at various stages of definition and clarity should be included. A concept such as "Gentleman" should be amenable to study just as "Vegetable" and "Furniture" have been studied. Conceptual structure for individuals may be assessed and the stimulus-category comparison process may be tested against norms provided by the individual. Collins and Quillian (1970) have cited idiosyncratic interpretations of category labels. The use of an individual's own category norms as a basis for assessing comparison processes can correct for such idiosyncratic interpretations.

The approach considered above concentrates on assessment of category structures in existence. There is an advantage to controlling the development of a category experientially. If a subject is required to learn a concept de novo, it is possible to control the relevant factors in its development. Hence, the evaluation of conceptual structure should be free of other unknown factors that might influence its development. As illustrated above the concept learning paradigm can be used for this purpose. It is also possible to control the evolution of a concept in a more naturalistic setting. For example, a concept not commonly known may be presented through discourse that includes examples. Manipulation of examples included is then hypothesized to determine the conceptual structure which can in turn be studied in the verification task. For example, concepts used may include ideas in the physical sciences or mathematics. The concept of a lattice may be presented to a subject, who possesses the appropriate naivete. Illustration of the concept with

exemplars should determine the conceptual structure. Concepts of a more subjective nature, such as "antihero," may be illustrated from literature, and its structure can be assessed. In this way acquisition of an abstract idea can be controlled in a manner reflecting ecological generality. Examination of structure may then reflect the process in a representative way.

Finally, the relationship of quantitative indices derived from the proposed system to fuzzy set inclusion functions derived from subjects responses to set inclusion relationships would be of great value in evaluating the adequacy of proposed quantitative measures, such as F, shown in (20). Hersh and Caramazza (1976) asked subjects to rate exemplars for degree of membership in category or goodness of exemplar in order to derive quantitative measures relevant to fuzzy logical operations. Such measures should relate monotonically to F derived for the same concept, if it is to reflect accurately the underlying process whereby fuzzy inclusion relationships are mediated. An inclusion function derived from subjects assessments for a given exemplar can be related to F_s, the corresponding fuzziness index for that comparison. A detailed study of the characteristics of underlying structure related to fuzzy inclusion functions can be used to modify quantitative indices such as F.

F. Relationship to Existing Data

The proposed approach should be reconcilable with existing evidence if it is to have any hope of describing human conceptual processes. The scheme presented involves both exemplar and feature comparison processes. Hence, its flexibility provides the possibility for post hoc explanations of results favoring either type of process. Such explanations would

prove unimpressive, and it could be argued that the ability to explain such a range of results proves to be a weakness of the approach. Nonetheless, it is necessary to see if there is any promise for this theoretical framework before considering more sensitive tests.

Smith, et al. (1974) presented a great deal of evidence that demonstrated that feature comparison processes can account for ease of comparison between stimuli and categories.

They showed that similarity to category, regardless of category size, accounts for such comparison data. The present formulation allows for a feature comparison process based upon an exemplar representation and hence could no doubt be used to account for such results. Other investigators (e.g., Wilkins, 1971 and Landauer & Meyer, 1972) have shown that reaction time for comparison of stimuli with categories vary with category size when representativeness of stimuli to category name is held constant. Such results suggest that a search of the set of elements of the category takes place in the comparison process. This evidence thus provides difficulty for the feature comparison point of view. Smith, et al. have criticized the research on methodological grounds. For example, one criticism of Wilkins' finding centered on the validity of his measure of category size, which was based upon normative data for production of category instances. However, it should be noted that any category size effect predicted by the present formulation would depend upon the size of the set of salient exemplars. Therefore, normative data of the sort used by Wilkins would be relevant to such a test. A detailed analysis of the category size issue is given by Smith, et al. and won't be considered here. However, insofar as category size effects

independent of similarity do exist, they can prove troublesome for a theory that relies solely on feature comparisons. In contrast, the proposed approach can accommodate an exemplar search process and account for such category size effects. Insofar as an initial search appears to be "holistic," as pointed out by Smith, et al., an initial search of salient exemplars seem promising. Related exemplars may be seen as collections of correlated features, in which case, an exemplar search would be reflected in a correlation of features used.

The concept identification results of Nahinsky, Oeschger, and Morgan, referred to in Section B, indicate that concept representation does involve sets of highly associated exemplars. Other concept identification research (e.g., Baum, 1954; Dattman & Israel, 1951; Heidbreder, 1946; and Katz & Denny, 1977) suggest that the nature of instances encountered en route to concept acquisition have an impact upon the development of the concept. Thus, it is not unreasonable to hypothesize that representation of concepts incorporate such contextual information. It is also important to note that the Nahinsky, Oeschger, and Morgan results suggested that featural resemblance of stimuli to instances frequently encountered during learning facilitated their classification. Thus, the evidence favored feature as well as exemplar matching processes, both of which are included in the proposed formulation.

Evidence deriving from another domain is supportive of representation in terms of salient exemplars. Tversky (1977) has pointed out the asymmetry involved in making similarity judgments, emphasizing, that such statements are generally asymmetrical. We say, "North Korea is like Red China " not, "Red China is like North Korea." The evidence Tversky

presented indicated that the choice of status for stimuli is determined by the relative salience of the stimuli involved in a similarity judgment. He presented a theory based upon feature comparisons between stimulus and referent, which need not concern us here. However, the fact that certain stimuli within a given category are more likely than others to serve as foci for similarity judgments suggests that representation of concepts may be, at least in part, on the basis of salient exemplars. It should be noted that these stimuli are specific examples, not composites as implied by the concept of prototype. Thus, typicality effects, as noted by Rosch (1975), for example, may reflect representation by highly associated stimuli rather than some "average" representative stimulus.

The system proposed introduces a new dimension to the study of category representation. The relationship among category members with respect to feature composition assumes importance in the representation system presented, whereas no other theory is explicitly concerned with this aspect of category structure. Thus, the investigations involving other approaches present no evidence necessary to distinguish them from the approach advanced herein. Certainly, there is no reason to expect that such evidence would be readily available.

It would be necessary to assess category structure for variables proposed here before evidence crucial to evaluate the described approach can be produced. Available data probably does not lend itself readily to the necessary quantification. However, it can be said that available evidence does not appear to be incompatible with proposed approach.

References

Barresi, J., Robbins, D., & Shain, K. Role of distinctive features in the abstraction of related concepts. Journal of Experimental Psychology: Human Learning and Memory, 1975, 1, 360-368.

Battig, W. F., & Montague, W. E. Category norms for verbal items in 56 categories: A replication and extension of the Connecticut Category Norms. Journal of Experimental Psychology Monograph, 1969, 80, (3, Pt. 2).

Baum, M. Simple concept learning as a function of intra-list generalization. Journal of Experimental Psychology, 1954, 47, 89-94.

Bolinger, D. The atomization of meaning. Language, 1965, 41, 555-573.

Bower, G. H., & Trabasso, T. R. Concept identification. In R. C. Atkinson (Ed.) Studies in Mathematical Psychology. Stanford: Stanford University Press, 1964.

Cohen, B. H., Bousfield, W. A., & Whitmarsh, G. A. Cultural norms for verbal items in 43 categories. (Office of Naval Research Tech. Rep. 22, Nonr 631 (00) Storrs: University of Connecticut, 1957.

Collins, A. M., & Quillian, M. R. Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior, 1969, 8, 240-248.

Dattman, P., & Israel, H. The order of dominance among conceptual capacities: An experimental test of Heidbreder's hypothesis. Journal of Psychology, 1951, 31, 147-160.

Falmagne, R. Construction of a hypothesis model for concept identification. Journal of Mathematical Psychology, 1970, 7, 60-96.

Garner, W. R., & Felfoldy, G. L. Integrality of stimulus dimensions in various types of information processing. Cognitive Psychology, 1970, 1, 225-241.

Greeno, L. W., & Simon, H. A. Process models and stochastic theories of simple concept formation. Journal of Mathematical Psychology, 1967, 4, 246-276.

Hayes Roth, B. Evolution of cognitive structures and processes. Psychological Review, 1977, 84, 260-278.

Hebb, D. O. Organization of Behavior. New York: Wiley, 1949.

Heidbreder, E. The attainment of concepts: I. Terminology and methodology. Journal of General Psychology, 1946, 35, 173-189.

Hersh, H. M., & Caramazza, A. A fuzzy set approach to modifiers and vagueness in natural language. Journal of Experimental Psychology: General, 1976, 105, 254-276.

Katz, A. N., & Denny, J. P. Memory-load and concreteness in the order of dominance effect for verbal concepts. Journal of Verbal Learning and Verbal Behavior, 1977, 16, 13-20.

Katz, J. J., & Fodor, J. A. The structure of a semantic theory. Language, 1963, 39, 170-210.

Kintsch, W. The representation of meaning in memory. Hillsdale, N. J.: L. Erlbaum Associates, 1974.

Labov, W. The boundaries of words and their meanings. In C.-J. N. Bailey & R. W. Shuy (Eds.), New ways of analyzing variation in English. Washington, D. C.: Georgetown University Press, 1973.

Lakoff, G. Hedges: A study in meaning criteria and the logic of fuzzy concepts. Journal of Philosophical Logic, 1973, 2, 458-508.

Landauer, T. K., & Freedman, J. L. Information retrieval from long-term memory: Category size and recognition time. Journal of Verbal Learning and Verbal Behavior, 1968, 7, 291-295.

Landauer, T. K., & Meyer, D. E. Category size and semantic memory retrieval. Journal of Verbal Learning and Verbal Behavior, 1972, 11, 539-549.

Levine, M. Hypothesis behavior by humans during discrimination learning. Journal of Experimental Psychology, 1966, 71, 331-338.

Loftus, E. F. Spreading activation within semantic categories: Comments on Rosch's "Cognitive representation of semantic categories." Journal of Experimental Psychology: General, 1975, 104, 234-240.

Meyer, D. E. On the representation and retrieval of stored semantic information. Cognitive Psychology, 1970, 1, 242-299.

Nahinsky, I. D. A hypothesis sampling model for conjunctive concept identification. Journal of Mathematical Psychology, 1970, 7, 293-316.

Nahinsky, I. D., Oeschger, D. E., & O'Leary, D. Stimulus memory and contextual cues in the abstraction process. Journal of Canadian Psychology, 1977, 31, 102-112.

Neumann, P. G. An attribute frequency model for the abstraction of prototypes. Memory & Cognition, 1974, 2, 241-248.

Neumann, P. G. Visual prototype formation with discontinuous representation of dimensions of variability. Memory & Cognition, 1977, 5, 187-197.

Oliver, R. R., & Hornsby, J. R. On equivalence. In J. R. Bruner, R. R. Oliver, & P. M. Greenfield (Eds.) Studies in Cognitive Growth, New York: Wiley, 1967.

Posner, M. I. Abstraction and the process of recognition. In G. H. Bower & J. T. Spence (Eds.), The psychology of learning and motivation (Vol. 3). New York: Academic Press, 1969.

Posner, M. I., & Keele, S. W. On the genesis of abstract ideas. Journal of Experimental Psychology, 1968, 77, 353-363.

Restle, F. The selection of strategies in cue learning. Psychological Review, 1962, 69, 329-343.

Rosch, E. Cognitive representation of semantic categories. Journal of Experimental Psychology: General, 1975, 104, 192-233.

Rosenberg, S., & Simon, H. A. Modeling semantic memory: Effects of presenting semantic information in different modalities. Cognitive Psychology, 1977, 9, 293-325.

Rumelhart, D. E., Lindsay, P. H., & Norman, D. A. A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), Organization and memory. New York: Academic Press, 1972.

Smith, E. E., Shoben, E. J., & Rips, L. J. Structure and Process in semantic memory: a featural model for semantic decision. Psychological Review, 1974, 81, 214-241.

Townsend, J. T. Some results concerning the identifiability of parallel and serial processes. British Journal of Mathematical and Statistical Psychology, 1972, 25, 168-199.

Tversky, A. Features of similarity. Psychological Review, 1977, 84, 327-352.

Wickens, T. D., & Millward, R. B. Attribute elimination strategies for concept identification with practiced subjects. Journal of Mathematical Psychology, 1971, 8, 453-480.

Wilkins, A. J. Conjoint frequency, category size, and categorization time. Journal of Verbal Learning and Verbal Behavior, 1971, 10, 382-385.

Wittgenstein, L. Philosophical Investigations. (Trans. by G. E. M. Anscombe) Oxford, England: Blackwell, 1953.

Zadeh, L. A. Fuzzy sets. Information and Control, 1965, 8, 338-353.

Zadeh, L. A. Outline of a new approach to the analysis of complex systems and decision processes. IEEE Transactions on Systems, Man, and Cybernetics, 1973, SMC-3, 28-44.